



Revision Guide for AMD NPT Family 0Fh Processors

Publication # 33610	Revision: 3.00
Issue Date: October 2006	

© 2006 Advanced Micro Devices, Inc. All rights reserved.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. The information contained herein may be of a preliminary or advance nature and is subject to change without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

AMD, the AMD Arrow logo, AMD Athlon, AMD Opteron, AMD Sempron, AMD Turion, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Revision History

Date	Revision	Description
October 2006	3.00	Initial public release.

Revision Guide for AMD NPT Family 0Fh Processors

The purpose of the *Revision Guide for AMD NPT Family 0Fh Processors* is to communicate updated product information to designers of computer systems and software developers. This revision guide includes information on the following products:

- AMD Athlon™ 64 processor
- AMD Athlon 64 X2 dual-core processor
- AMD Athlon 64 FX dual-core processor
- Dual-Core AMD Opteron™ processor
- AMD Turion™ 64 X2 Mobile Technology
- AMD Sempron™ processor

This guide consists of three major sections:

- **Processor Identification:** This section, starting on page 6, shows how to determine the processor revision, program and display the processor name string, and construct the processor name string.
- **Product Errata:** This section, starting on page 13, provides a detailed description of product errata, including potential effects on system operation and suggested workarounds. An erratum is defined as a deviation from the product's specification, and as such may cause the behavior of the processor to deviate from the published specifications.
- **Documentation Support:** This section, starting on page 57, provides a listing of available technical support resources.

Revision Guide Policy

Occasionally, AMD identifies product errata that cause the processor to deviate from published specifications. Descriptions of identified product errata are designed to assist system and software designers in using the processors described in this revision guide. This revision guide may be updated periodically.

Processor Identification

This section shows how to determine the processor revision, program and display the processor name string, and construct the processor name string.

Revision Determination

Figure 1 shows the format of the value returned in EAX by CPUID Function 1.




Figure 1. Format of CPUID Value Returned by Function 1

Table 1 through 4 shows the identification number returned by the CUID instruction for each revision of the processor.

Table 1. CUID Values for Revisions of AMD Opteron™ Processors

CUID Function 1 EAX Value	
Rev.	Dual-Core AMD Opteron™ Processor
JH-F2	Socket F (1207) 00040F12h Socket AM2 00040F32h

Table 2. CUID Values for Revisions of AMD Athlon™ 64 Processors

CUID Function 1 EAX Value			
Rev.	AMD Athlon™ 64 Processor	AMD Athlon™ 64 X2 Dual-Core Processor	AMD Athlon™ 64 FX Dual-Core Processor
JH-F2	N/A	Socket AM2 00040F32h	Socket AM2 00040F32h
BH-F2	N/A	Socket AM2 00040FB2h	N/A
DH-F2	Socket AM2 00040FF2h Socket AM2 00050FF2h	N/A	N/A

Table 3. CUID Values for Revisions of AMD Sempron™ Processors

CUID Function 1 EAX Value	
Rev.	AMD Sempron™ Processor
DH-F2	Socket AM2 00040FF2h Socket AM2 00050FF2h

Table 4. CUID Values for Revisions of AMD Turion™ 64 Processors

CUID Function 1 EAX Value	
Rev.	AMD Turion™ 64 X2 Mobile Technology
JH-F2	N/A
BH-F2	Socket S1g1 00040F82h
DH-F2	N/A

Programming and Displaying the Processor Name String

This section, intended for BIOS programmers, describes how to program and display the 48-character processor name string that is returned by CUID Fn8000_000[4:2]. The hardware or cold reset value

of the processor name string is 48 ASCII NUL characters, so the BIOS must program the processor name string before any general purpose application or operating system software uses the extended functions that read the name string. It is common practice for the BIOS to display the processor name string and model number whenever it displays processor information during boot up.

Note: *Motherboards that do not program the proper processor name string and model number will not pass AMD validation and will not be posted on the AMD Recommended Motherboard Website.*

The name string must be ASCII NUL terminated and the 48-character maximum includes that NUL character.

The processor name string is programmed by MSR writes to the six MSR addresses covered by the range C001_00[35:30]h. Refer to the *BIOS and Kernel Developer's Guide for NPT Processors*, order# 32559, for the format of how the 48-character processor name string maps to the 48 bytes contained in the six 64-bit registers of MSR C001_00[35:30]h.

The processor name string is read by CPUID reads to a range of CPUID functions covered by CPUID Fn8000_000[4:2]. Refer to CPUID Fn8000_000[4:2] in the *BIOS and Kernel Developer's Guide for NPT Processors*, order# 32559, for the 48-character processor name string mapping to the 48 bytes contained in the twelve 32-bit registers of CPUID Fn8000_000[4:2].

Constructing the Processor Name String

This section describes how to construct the processor name string. BIOS uses the following fields to create the name string:

- **BrandId[15:0]** is from CPUID Fn8000_0001_EBX.
- **BaseModel[1:0]** is from CPUID Fn8000_0001_EAX.
 - **PwrLmt[3:0]** is defined to be BrandID[8:6,14]. See the *BIOS and Kernel Developer's Guide for NPT Processors*, order# 32559, for details on BIOS use of PwrLmt[3:0].
 - **BrandTableIndex[4:0]** is defined to be BrandID[13:9]. This field is used in Table 6 on page 10 through Table 8 on page 12 to index into the processor name string tables.
 - **NN[6:0]** is defined to be BrandID[15,5:0]. This field is used to create the model number in the name string. Use of NN[6:0] is explained in Table 9 on page 12.
- **Socket[1:0]** is defined to be BaseModel[1:0]. 00b = Socket S1g1, 01b = Socket F (1207), 10b = Reserved, 11b = Socket AM2. This field is used to index the appropriate name string table from Table 5.
- **CmpCap[1:0]:** CMP Capable. Specifies the number of processor cores enabled on the device. 00b=1 processor core; 01b=2 processor cores; 10b=Reserved; 11b=Reserved. CmpCap[1:0] is located in the Northbridge Capabilities Register (Dev:3xE8[13:12]) and is defined in the *BIOS and Kernel Developer's Guide for NPT Processors*, order# 32559. This field is used to index the appropriate name string in Table 6 through 8.

The name string is formed by selecting the appropriate name string table from Table 5 using the socket type, and then indexing into that table with CmpCap[0], BrandTableIndex[4:0], and PwrLmt[3:0]. The model number is calculated using NN[6:0] and the equation from Table 9 on page 12 as indexed from the note for each name string specified in Table 6 on page 10 through Table 8 on page 12. If NN[6:0]=0, then the name string is “AMD Engineering Sample”.

Table 5. Using the Socket Type to Select the Name String Table

Socket	Name String Table
F (1207)	See Table 6 on page 10
AM2	See Table 7 on page 11.
S1g1	See Table 8 on page 12.

Table 6. Processor Name String Table for Socket F (1207)

CmpCap [0]	Brand Table Index [4:0]	PwrLmt [3:0]	Name String	Notes	Description
-	00h	0h	AMD Engineering Sample	-	Reserved
1b (Dual Core)	01h	6h	Dual-Core AMD Opteron(tm) Processor 22RR HE	1	DP Server Dual-Core
	01h	Ah	Dual-Core AMD Opteron(tm) Processor 22RR	1	DP Server Dual-Core
	01h	Ch	Dual-Core AMD Opteron(tm) Processor 22RR SE	1	DP Server Dual-Core
	04h	6h	Dual-Core AMD Opteron(tm) Processor 82RR HE	1	MP Server Dual-Core
	04h	Ah	Dual-Core AMD Opteron(tm) Processor 82RR	1	MP Server Dual-Core
	04h	Ch	Dual-Core AMD Opteron(tm) Processor 82RR SE	1	MP Server Dual-Core
All other codes			AMD Processor model unknown	-	This case can occur if the processor is upgraded but not the BIOS.

Table 7. Processor Name String Table for Socket AM2

CmpCap [0]	Brand Table Index [4:0]	PwrLmt [3:0]	Name String	Notes	Description
-	00h	0h	AMD Engineering Sample	-	Reserved
0b (Single Core)	04h	4h	AMD Athlon(tm) 64 Processor TT00+	3	Desktop/DTR Client
	04h	8h	AMD Athlon(tm) 64 Processor TT00+	3	Desktop/DTR Client
	06h	4h	AMD Sempron(tm) Processor TT00+	3	Desktop/DTR Client
	06h	8h	AMD Sempron(tm) Processor TT00+	3	Desktop/DTR Client
1b (Dual Core)	01h	Ah	Dual-Core AMD Opteron(tm) Processor 12RR	1	UP Server Dual-Core
	01h	Ch	Dual-Core AMD Opteron(tm) Processor 12RR SE	1	UP Server Dual-Core
	04h	2h	AMD Athlon(tm) 64 X2 Dual Core Processor TT00+	3	Dual-Core Client
	04h	6h	AMD Athlon(tm) 64 X2 Dual Core Processor TT00+	3	Dual-Core Client
	04h	8h	AMD Athlon(tm) 64 X2 Dual Core Processor TT00+	3	Dual-Core Client
	05h	Ch	AMD Athlon(tm) 64 FX-ZZ Dual Core Processor	4	Desktop Client Dual-Core
All other codes			AMD Processor model unknown	-	This case can occur if the processor is upgraded but not the BIOS.

Table 8. Processor Name String Table for Socket S1g1

CmpCap [0]	Brand Table Index [4:0]	PwrLmt [3:0]	Name String	Notes	Description
-	00h	0h	AMD Engineering Sample	-	Reserved
1b (Dual Core)	02h	Ch	AMD Turion(tm) 64 X2 Mobile Technology TL-YY	5	UP Client Dual-Core
All other codes			AMD Processor model unknown	-	This case can occur if the processor is upgraded but not the BIOS.

Table 9. Model Number Calculation

Number	Note
1	RR = -1 + NN. For example, 000011b represents “2”, 100001b represents “32”. Values 000001b (1) to 000010b (2), and 100010b (34) to 111111b (63) are reserved.
2	PP = 26 + NN. For example, 000001b represents “27”, 111111b represents “89”.
3	TT = 15 + CmpCap*10 + NN. For example, if CmpCap = 00b, 000001b represents “16”, 111111b represents “78”. If CmpCap = 01b, 000001b represents “26”, 111111b represents “88”.
4	ZZ = 57 + NN. For example, 000001b represents “58”, 100001b represents “90”. Values 100010b (34) to 111111b (63) are reserved.
5	YY = 29 + NN. For example, 000001b represents “30”, 111111b represents “92”.
Note: All of the model codes are expressed in decimal. NN is defined on page 8. If NN=0 then the name string is “AMD Engineering Sample”.	

Product Errata

This section documents product errata for the processors. A unique tracking number for each erratum has been assigned within this document for user convenience in tracking the errata within specific revision levels. Table 10 cross-references the revisions of the part to each erratum. An “X” indicates that the erratum applies to the revision. The absence of an “X” indicates that the erratum does not apply to the revision. An “*” indicates advance information that the erratum has been fixed but not yet verified. “No fix planned” indicates that no fix is planned for current or future revisions of the processor.

Note: *There may be missing errata numbers. Errata that have been resolved from early revisions of the processor have been deleted, and errata that have been reconsidered may have been deleted or renumbered.*

Table 10. Cross-Reference of Product Revision to Errata

No.	Errata Description	Revision Number		
		JH-F2	BH-F2	DH-F2
1	Inconsistent Global Page Mappings Can Lead to Machine Check Error	No fix planned		
57	Some Data Cache Tag Eviction Errors Are Reported As Snoop Errors	No fix planned		
60	Single Machine Check Error May Report Overflow	No fix planned		
75	APIC Timer Accuracy Across Power Management Events	No fix planned		
77	Long Mode CALLF or JMPF May Fail To Signal GP When Callgate Descriptor is Beyond GDT/LDT Limit	No fix planned		
89	Potential Deadlock With Locked Transactions	No fix planned		
112	Self-Modifying Code May Execute Stale Instructions	X	X	X
122	TLB Flush Filter May Cause Coherency Problem in Multiprocessor Systems	No fix planned		
125	Processor Temporarily Stops Issuing Refresh Requests While Exiting S3 State	No fix planned		
126	Incorrect Output Drive Strength in Systems Using Both x4 and Non-x4 RDIMMs	X	X	X
130	Interrupt Shadow Property May Not Be Observed In SVM Mode	X	X	X
131	Systems May Deadlock Waiting for a Probe Response	No fix planned		
133	Internal Termination Missing on Some Test Pins	X	X	X
138	AltVidTriEn Bit Causes System Hang	X	X	X
140	TSC Offset Is Not Added to Reads of the TSC MSR 0x10 When Running a Virtualized Guest	X	X	X
141	Inaccurate Internal Triggering of HTC/STC Feature	X	X	X
142	DQS Receiver Enable Delays Incorrectly Set In Mismatched DIMM Mode	No fix planned		
144	CLFLUSH to Shadow RAM Address Will Not Invalidate	No fix planned		
149	Clock Jitter on MEMCLK Pins During Write Transactions	No fix planned		
151	Incorrect Programming of the Specific End of Interrupt (SEOI) Register Results in Processor Hang	No fix planned		

Table 10. Cross-Reference of Product Revision to Errata (Continued)

No.	Errata Description	Revision Number		
		JH-F2	BH-F2	DH-F2
152	DRAM Controller Violates tXRSd Requirement when Exiting Self Refresh Mode After Some P-State Changes	No fix planned		
153	Potential System Hang in Multiprocessor Systems With ≥14 Cores	X		
156	Read Request Incorrectly Generated for Invalid GART Page Table Entries	No fix planned		
158	HLT Initiated C1 State Transitions May Lead to Processor Hang In SVM Guest Mode	X	X	X
159	Clock Jitter on MEMCLK Pins During Read Transactions			X
161	Performance-Monitoring Counters Do Not Count Code Address Matches	X	X	X
162	Writes to Read-Only APIC Register Cause Processor Hang	No fix planned		
164	DRAM Refresh Controller Not Enabled After Using BIOS Controlled DRAM Initialization	X	X	X
165	#VMEXIT(INVALID) Unconditionally Clears EVENTINJ Field In VMCB	X	X	X
166	FXSAVE/FXRSTOR Instructions Use 64-bit Format in Compatibility Mode	X	X	X
167	System Hang with Chipset Initiated PROCHOT_L Throttling When Clock Divisor is Greater Than 64	X	X	X
168	System May Hang When Exiting C1E or C3	X	X	X
169	System May Hang Due to DMA or Stalled Probe Response	No fix planned		
170	In SVM Mode Incorrect Code Bytes May Be Fetched After A World Switch	X	X	X
171	Instruction Break Point On VMRUN Instruction Leads To Unpredictable System Behavior	No fix planned		
172	Some Registered DIMMs Incompatible With Address Parity Feature	X		
181	Asserting LDTSTOP_L Before DRAM is Initialized May Cause System Hang	X	X	X

Table 11 cross-references the errata to each processor segment. An empty cell signifies that the erratum does not apply to the processor segment. “X” signifies that the erratum applies to the processor segment. “N/A” signifies that the erratum does not apply to the processor segment due to the silicon revision.

Table 11. Cross-Reference of Errata to Processor Segments

Errata Number	Dual-Core AMD Opteron™ Processor (Socket AM2)	Dual-Core AMD Opteron™ Processor (Socket F 1207)	AMD Athlon™ 64 Processor	AMD Athlon™ 64 X2 Dual-Core Processor	AMD Athlon™ 64 FX Dual-Core Processor	AMD Sempron™ Processor	AMD Turion™ 64 X2 Mobile Technology
1	X	X	X	X	X	X	X
57	X	X	X	X	X	X	X
60	X	X	X	X	X	X	X
75	X	X	X	X	X	X	X
77	X	X	X	X	X	X	X
89	X	X	X	X	X	X	X
112	X	X	X	X	X	X	X
122	X	X		X	X		X
125	X	X	X	X	X	X	X
126		X					
130	X	X	X	X	X	X	X
131	X	X	X	X	X	X	X
133	X	X	X	X	X	X	X
138							X
140	X	X	X	X	X	X	X
141	X	X	X	X	X	X	X
142			X	X	X	X	X
144	X	X	X	X	X	X	X
149	X	X	X	X	X	X	X
151	X	X	X	X	X	X	X
152	X	X	X	X	X	X	X
153		X					
156	X	X	X	X	X	X	X
158	X	X	X	X	X	X	X
159			X			X	
161	X	X	X	X	X	X	X
162	X	X	X	X	X	X	X
164	X	X	X	X	X	X	X
165	X	X	X	X	X	X	X
166	X	X	X	X	X	X	X
167							X
168							X

Table 11. Cross-Reference of Errata to Processor Segments (Continued)

Errata Number	Dual-Core AMD Opteron™ Processor (Socket AM2)	Dual-Core AMD Opteron™ Processor (Socket F 1207)	AMD Athlon™ 64 Processor	AMD Athlon™ 64 X2 Dual-Core Processor	AMD Athlon™ 64 FX Dual-Core Processor	AMD Sempron™ Processor	AMD Turion™ 64 X2 Mobile Technology
169	X	X	X	X	X	X	X
170	X	X	X	X	X	X	X
171	X	X	X	X	X	X	X
172		X					
181	X	X	X	X	X	X	X

1 Inconsistent Global Page Mappings Can Lead to Machine Check Error

Description

If the same linear to physical mapping exists in multiple CR3 contexts, and that mapping is marked global in one context and not global in another context, then a machine check error may be reported by the TLB error detection logic (depending on the specific access pattern and TLB replacements encountered).

Potential Effect on System

In the somewhat unlikely event that all required conditions are present (including the effects of the TLB replacement policy), then an unexpected machine check error may be reported. If the erratum occurs in the instruction cache TLB (L1 or L2), the apparent error is logged and corrected. If the erratum occurs in the data cache TLB (L1 or L2), the apparent error is logged and reported as an uncorrectable machine check error.

Suggested Workaround

None required. This is not expected to occur in real systems.

Fix Planned

No

57 Some Data Cache Tag Eviction Errors Are Reported As Snoop Errors

Description

In some cases, the machine check error code on a data cache (DC) tag array parity error erroneously classifies an eviction error as a snoop error.

The common cases of cache line replacements and external probes are classified correctly (as eviction and snoop respectively). The erroneous cases occur when a tag error is detected during a DC eviction that was generated by a hardware prefetch, a cache line state change operation, or a number of other internal microarchitectural events. In such cases, the error code logged in the DC Machine Check Status register (MC0_STATUS, MSR 0x401) erroneously indicates a snoop error.

Potential Effect on System

Internally detected DC tag errors may be reported to software as having been detected by snoops. Depending upon machine check software architecture, the system response to such errors may be broader than necessary.

Suggested Workaround

None required.

Fix Planned

No

60 Single Machine Check Error May Report Overflow

Description

A single parity error encountered in the data cache tag array may incorrectly report the detection of multiple errors, as indicated by the overflow bit of the DC Machine Check Status register (bit 62 of MSR 0x401).

Potential Effect on System

System software may be informed of a machine check overflow when only a single error was actually encountered.

Suggested Workaround

Do not rely on the state of the OVER bit in the DC Machine Check Status register.

Fix Planned

No

75 APIC Timer Accuracy Across Power Management Events

Description

The APIC timer may be inaccurate by up to 1 μ s across each use of S1 or LDTSTOP_L initiated HyperTransport™ link width/frequency changes.

Potential Effect on System

No observable system impact expected.

Suggested Workaround

None.

Fix Planned

No

77 Long Mode CALLF or JMPF May Fail To Signal GP When Callgate Descriptor is Beyond GDT/LDT Limit

Description

If the target selector of a far call or far jump (CALLF or JMPF) instruction references a 16-byte long mode system descriptor where any of the last 8 bytes are beyond the GDT or LDT limit, the processor fails to report a General Protection fault.

Potential Effect on System

None expected, since the operating system typically aligns the GDT/LDT limit such that all descriptors are legal. However, in the case of erroneous operating system code, the above described GP fault will not be signaled, resulting in unpredictable system failure.

Suggested Workaround

None required, it is anticipated that long mode operating system code will ensure the GDT and LDT limits are set high enough to cover the larger (16-byte) long mode system descriptors.

Fix Planned

No

89 Potential Deadlock With Locked Transactions

Description

Downstream non-posted requests to devices that are dependent on the completion of an upstream non-posted request can cause a deadlock in the presence of transactions resulting in bus locks, as shown in the following two scenarios:

1. A downstream non-posted read to the LPC bus occurs while an LPC bus DMA is in progress. The legacy LPC DMA blocks downstream traffic until it completes its upstream reads.
2. A downstream non-posted read is sent to a device that must first send an upstream non-posted read before it can complete the downstream read.

In both cases, a locked transaction causes the upstream channel to be blocked, causing the deadlock condition.

Potential Effect on System

The system fails due to a bus deadlock.

Suggested Workaround

BIOS should set the DisIOReqLock bit (bit 3 in NB_CFG, MSR C001_001F).

Fix Planned

No

112 Self-Modifying Code May Execute Stale Instructions

Description

The processor may execute stale instructions in a situation involving a store instruction modifying a younger instruction within two cache lines of each other and an internal processor trap condition occurring in a small window:

1. The fetch for the store and the younger instruction must hit in the instruction cache (IC) and are brought into the processor pipeline.
2. The store speculatively executes, and prior to invalidating the line in the IC, an internal trap event occurs.
3. This internal trap event must redirect the instruction fetch to a point in the instruction stream just before the store.
4. The fetch associated with the redirect must hit in the IC before the first execution of the store (in step 2) has invalidated the line containing the target.
5. During the small window of time between the refetch of the younger instruction and the fetch of the last instruction in that cache line from the IC, the speculative store from step 2 invalidates the IC line which contains the younger instruction.

Potential Effect on System

In the unlikely event that the above conditions occur, the processor will execute stale instruction(s).

Just in time (JIT) compilers lack the proximity of the store instruction to the modified code and thus are not affected.

Suggested Workaround

This scenario was contrived in a highly randomized simulation environment and is not expected to occur in a real system. In the unlikely event that the erratum is observed, a serializing instruction can be inserted prior to executing the modified code.

Fix Planned

Yes

122 TLB Flush Filter May Cause Coherency Problem in Multiprocessor Systems

Description

Under highly specific internal timing conditions in a multiprocessor configuration, coherency problems may arise between the page tables in memory and the translations stored in the on-chip TLBs. This can result in the possible use of stale translations even after software has performed a TLB flush.

Potential Effect on System

Unpredictable system failure. This scenario has only been observed in a highly randomized synthetic stress test.

Suggested Workaround

In multiprocessor or multicore systems, disable the TLB flush filter by setting HWCR.FFDIS (bit 6 of MSR 0xC001_0015).

Fix Planned

No

125 Processor Temporarily Stops Issuing Refresh Requests While Exiting S3 State

Description

While exiting low power state S3 (Suspend to RAM state), the processor brings the DRAMs out of self refresh mode and waits 200 μ s to allow the internal circuitry to stabilize (e.g. DLLs to lock etc.). During this period, the DRAM controller will not issue refresh requests to the DRAMs for 200 μ s.

Potential Effect on System

The DRAM may potentially lose contents while exiting S3 state, resulting in unexpected system behavior.

Suggested Workaround

The BIOS should perform the following steps in the given order when exiting from S3 (Suspend to RAM) state in order to resume the refresh requests to DRAM immediately:

1. Restore memory controller registers as normal.
2. Set the EnDramInit bit (Dev:2x7C[31]), clear all other bits in the same register).
3. Wait 750 μ s.
4. Clear the EnDramInit bit.
5. Read the value of Dev:2x80 and write that value back to Dev:2x80.
6. Set the exit from the self refresh bit (Dev:2x90[1]).
7. Clear the exit from self refresh bit immediately.

Note: Steps 6 and 7 must be executed in a single 64-byte aligned uninterrupted instruction stream.

Fix Planned

No

126 Incorrect Output Drive Strength in Systems Using Both x4 and Non-x4 RDIMMs

Description

The processor provides the capability of setting different drive strengths for certain groups of DRAM outputs, including the MA_DQS[16:9] and MB_DQS[16:9] outputs. The drive strengths are programmed via the Output Driver Compensation Control registers (Dev:2x9C; Indexes 00h and 20h). If the DataDrvStren and DqsDrvStren fields in these registers are not programmed to the same strength values, the output drivers listed above may have the incorrect drive strength applied in systems that use both x4 and non-x4 registered DIMMs.

Potential Effect on System

The signal integrity of the DQS outputs may be affected negatively. This may limit the DRAM frequencies and/or the number of DIMMs supported by the processor.

Suggested Workaround

Do not program the DataDrvStren and DqsDrvStren fields to different settings in system configurations using both x4 and non-x4 registered DIMMs.

Fix Planned

Yes

130 Interrupt Shadow Property May Not Be Observed In SVM Mode

Description

In guest mode, if the STI instruction takes a debug exception due to a single step and IF=0 and the BP# exception is intercepted by the hypervisor, then the pending INTR will be taken before executing the shadowed instruction.

Potential Effect on System

If the guest is unable to tolerate INTRs after the above scenario, unpredictable behavior may occur.

Suggested Workaround

The hypervisor should not single step the guest across STI instructions with IF=0.

Fix Planned

Yes

131 Systems May Deadlock Waiting for a Probe Response

Description

Under a highly specific and detailed set of internal timing conditions, the Northbridge System Request Queue (SRQ) may stall a probe response leading to a deadlock.

Potential Effect on System

Deadlock or a machine check error due to a watchdog timer time-out.

Suggested Workaround

BIOS should set NB_CFG Register[20] (MSR C001_001Fh). No loss of performance results from this workaround.

Systems implementing the workaround for erratum 169 should not apply this workaround.

Fix Planned

No

133 Internal Termination Missing on Some Test Pins

Description

Internal termination resistors are missing on certain test pins.

Potential Effect on System

Floating inputs on test pins resulting in unpredictable operation.

Suggested Workaround

Termination must be added to the motherboard as shown in the following table.

Pin Name	Termination Value	Voltage Rail	Socket F (1207) Pin No.	Socket AM2 Pin No.	Socket S1g1 Pin No.
TEST18	300 ohms	VSS	F23	E9	H10
TEST19	300 ohms	VSS	F20	F10	G9
TEST21	300 ohms	VSS	AJ20	AL8	AB8
TEST26	300 ohms	VDDIO	AF18	AK5	AE6

Fix Planned

Yes

138 AltVidTriEn Bit Causes System Hang

Description

Setting the AltVidTriEn bit (Dev:3xD8[26]) correctly tristates the HyperTransport technology clocks when entering the C3 state with AltVid enabled, but upon exiting C3, clocks are enabled only after the VID transitions back to normal operating voltage. The HyperTransport technology specification requires clocks to be enabled within 400 ns of LDTSTOP_L deassertion but in this case clocks are enabled after 10 us.

Potential Effect on System

The system hangs.

Suggested Workaround

The BIOS should not set the AltVidTriEn bit (Dev:3xD8[26]). Disabling this feature results in a negligible power increase.

Fix Planned

Yes

140 TSC Offset Is Not Added to Reads of the TSC MSR 0x10 When Running a Virtualized Guest

Description

When a guest is running, the execution of the RDTSC or RDTSCP instructions returns the value of the architectural time stamp counter's contents plus the offset in the guest's VMCB. However, if the guest reads from MSR 0x10, the value returned is the architectural time stamp value without the offset added.

Potential Effect on System

A program running in a virtualized guest that reads MSR 0x10 receives a different value than it gets from RDTSC or RDTSCP.

Suggested Workaround

A hypervisor should intercept guest reads of MSR 0x10 by setting the corresponding bit in the MSR permissions map and manually add the TSC offset.

Fix Planned

Yes

141 Inaccurate Internal Triggering of HTC/STC Feature

Description

The internal thermal sensor used for the hardware thermal control (HTC) and software thermal control (STC) features is inaccurate.

Potential Effect on System

Activation of HTC/STC is inconsistent and does not provide reliable thermal protection.

Suggested Workaround

BIOS should not enable the HTC/STC features. Systems should be designed with conventional thermal control/throttling methods or utilize PROCHOT_L functionality based on temperature measurements from an analog thermal diode (THERMDA/THERMDC).

Fix Planned

Yes

142 DQS Receiver Enable Delays Incorrectly Set In Mismatched DIMM Mode

Description

The BIOS programmable DQS receiver delays in the DRAM controller interface use incorrect values when operating in mismatched DIMM mode.

Potential Effect on System

DRAM failures resulting in unpredictable system behavior.

Suggested Workaround

When operating in mismatched DIMM mode, BIOS should program identical values in the following DQS Receiver Enable timing registers:

- (Dev:2x9C) Offsets 30h and 36h should have identical DqsRcvEn delay values.
- (Dev:2x9C) Offsets 33h and 39h should have identical DqsRcvEn delay values.

Fix Planned

No

144 CLFLUSH to Shadow RAM Address Will Not Invalidate

Description

WrDram and RdDram bits in extended MTRR type registers are used to copy BIOS ROM to corresponding DRAM and then execute out of DRAM. When these are configured to direct writes to ROM (WrMem =0b) and reads to DRAM (RdMem =01b), the CLFLUSH instruction will not invalidate shadow RAM addresses in the cache.

Potential Effect on System

CLFLUSH instruction will be ineffective for shadow RAM space.

Suggested Workaround

Use the WBINVD instruction instead of CLFLUSH in shadow RAM space.

Fix Planned

No

149 Clock Jitter on MEMCLK Pins During Write Transactions

Description

Under certain conditions excessive clock jitter (tJIT(per) and tJIT(cc)) is observed on the MEMCLK pins during DDR write transactions above DDR533 (includes M[B, A][1:0]_CLK_H/L[2:0] pins on AM2 package, M[B, A]0_CLK_H/L[2:1] pins on S1g1 package, and M[B, A][3:0]_CLK_H/L pins on socket F (1207) package). The amount of jitter on the MEMCLK pins depends on the data pattern being driven during DRAM writes.

Potential Effect on System

None expected. No functional failures have been observed due to this issue.

Suggested Workaround

None required.

Fix Planned

No

151 Incorrect Programming of the Specific End of Interrupt (SEOI) Register Results in Processor Hang

Description

A hypervisor (or other software) that incorrectly writes to the Extended APIC Specific End of Interrupt (SEOI) register (APIC offset 420h) before enabling SEOI generation causes the processor to hang.

Potential Effect on System

The system hangs.

Suggested Workaround

Prior to writing to the Extended APIC SEOI register, hypervisors (or other software) should enable SEOI generation by setting SeoiEn in the Extended APIC Control register (APIC offset 410h, bit 1).

Fix Planned

No

152 DRAM Controller Violates tXSRD Requirement when Exiting Self Refresh Mode After Some P-State Changes

Description

On a P-state change in which the FID increases, the DRAM controller may issue a read prior to 200 MEMCLKs after CKE transitions high when exiting self refresh mode. This is a violation of the tXSRD requirement of the DDR2 SDRAM specification.

Potential Effect on System

The system may hang or the memory read may return invalid data.

Suggested Workaround

BIOS should set the Address Timing Control Register[CkeSetup] (function 2, offset 9C, index 04h and 24h, bit 5) and the Address Timing Control Register[CkeFineDelay] (function 2, offset 9C, index 04h and 24h, bits 4:0) as shown in the following tables.

SO-DIMM for Socket S1g1 Processors

DRAM Speed	CKE Setup	CKE Fine Delay
DDR2-400	0b	0h
DDR2-533	0b	0h
DDR2-667	0b	0h

Unbuffered DIMM for Socket AM2 Processors

DRAM Speed	CKE Setup	CKE Fine Delay
DDR2-400	0b	0h
DDR2-533	0b	0h
DDR2-667	1b	0h
DDR2-800	1b	0h

Registered DIMM for Socket F (1207) Processors

DRAM Speed	CKE Setup	CKE Fine Delay
DDR2-400	0b	0h
DDR2-533	0b	0h
DDR2-667	0b	0h

Fix Planned

No

153 Potential System Hang in Multiprocessor Systems With ≥ 14 Cores

Description

Under a highly specific and detailed set of internal timing conditions the crossbar flow control buffer pointer can become corrupted, causing incorrect information to be captured and resulting in a system hang.

Potential Effect on System

The system hangs.

Suggested Workaround

A BIOS workaround can be implemented to modify the buffer allocation scheme to limit the maximum number of crossbar command buffers to 49, with some performance impact. The AMD recommended buffer allocation for ladder and twisted ladder configurations are listed in the tables below. This requires modifications to the configuration registers Dev:0x90h, B0h, D0h, Dev:3x70h, and Dev:3x78h (command buffers only) and assumes default values in Dev:3x7Ch.

Outer Node Virtual Channel Command Buffer Allocation

Link	Request	Posted Request	Response	Probe	Number of Command Buffers
Coherent links	1	1	6	4	12
Non-coherent links	5	4	1	0	10
SRI	2	3	3	0	8
MCT	0	0	5	2	7

Total Command Buffer Allocation Per Outer Node

Allocation	Outer Node
12	Coherent link
12	Coherent link
10	Non-coherent link
8	SRI (includes three buffers from FreeList Buffer Count register)
7	MCT
49	Total

Dev:0x90h, B0h, D0h LDT Buffer Count Register Settings for Outer Nodes

Mnemonic	Req	PReq	Rsp	Probe
Register Bits	[3:0]	[7:4]	[11:8]	[15:12]
Coherent Links	1	1	6	4
Non-Coherent Links	5	4	1	0

Dev:3x70h SRI to XBAR Buffer Count Register Settings for Outer Nodes

Mnemonic	DReq	DPReq	UReq	UPReq	URsp
Register Bits	[29:28]	[31:30]	[1:0]	[5:4]	[9:8]
Value	1	1	1	1	1

Dev:3x78h MCT to XBAR Buffer Count Register Settings for Outer Nodes

Mnemonic	Rsp	Prb
Register Bits	[11:8]	[14:12]
Value	5	2

Inner Node Virtual Channel Command Buffer Allocation

Link	Request	Posted Request	Response	Probe	Number of Command Buffers
Coherent links	2	1	5	4	12
SRI	1	2	3	0	6
MCT	0	0	5	2	7

Total Command Buffer Allocation Per Inner Node

Allocation	Inner Node
12	Coherent link
12	Coherent link
12	Coherent link
6	SRI (includes three buffers from FreeList Buffer Count register)
7	MCT
49	Total

Dev:0x90h, B0h, D0h LDT Buffer Count Register Settings for Inner Nodes

Mnemonic	Req	Preq	Rsp	Probe
Register Bits	[3:0]	[7:4]	[11:8]	[15:12]
Coherent links	2	1	5	4

Dev:3x70h SRI to XBAR Buffer Count Register Settings for Inner Nodes

Mnemonic	DReq	DPReq	UReq	UPReq	URsp
Register Bits	[29:28]	[31:30]	[1:0]	[5:4]	[9:8]
Value	0	0	1	1	1

Dev:3x78h MCT to XBAR Buffer Count Register Settings for Inner Nodes

Mnemonic	Rsp	Prb
Register Bits	[11:8]	[14:12]
Value	5	2

Fix Planned

Yes

156 Read Request Incorrectly Generated for Invalid GART Page Table Entries

Description

When a read request is made to an invalid address within the GART aperture (i.e., when the valid bit is not set in the GART page table entry), the Northbridge incorrectly generates a read request to the page specified in the GART page table entry (PTE).

Potential Effect on System

A read request to an arbitrary address may result in undefined behavior.

Suggested Workaround

Operating system software should map a valid guard page after each contiguous region of valid pages in the aperture. Operating system software should also ensure that the physical page address of all invalid (unused) GART entries points to a guard page. The guard page should point to a valid physical page in DRAM; the same physical guard page can be used to guard all allocations and entries. The placement of the guard page is operating system dependent, but one solution is to assign a value of 0x0 for the physical page address of the guard page.

Fix Planned

No

158 HLT Initiated C1 State Transitions May Lead to Processor Hang In SVM Guest Mode

Description

If C1 clock ramping is enabled in SVM guest context, and if HLT is executed immediately after STI or as a first guest instruction following a VMRUN with interrupt shadow enabled, a pending virtual interrupt will lead to processor hang.

Potential Effect on System

C1 state transitions in SVM guest mode may lead to processor hang.

Suggested Workaround

Hypervisors should not execute the VMRUN instruction with V_IRQ = 1 and HLT Intercept = 0 or they should disable C1 clock ramping in guest mode.

Fix Planned

Yes

159 Clock Jitter on MEMCLK Pins During Read Transactions

Description

Under certain conditions excessive clock jitter, tJIT(cc), is observed on the MEMCLK pins during DDR read transactions above DDR533 (includes M[B, A][1:0]_CLK_H/L[2:0] pins on socket AM2 processors, and M[B, A]0_CLK_H/L[2:1] pins on socket S1g1 processors). The amount of jitter on the MEMCLK pins depends on the data pattern being driven during DRAM reads.

Potential Effect on System

None expected. No functional failures have been observed due to this issue.

Suggested Workaround

None required.

Fix Planned

No

161 Performance-Monitoring Counters Do Not Count Code Address Matches

Description

The performance-monitoring counters do not count code address matches with DR0-3 unless the corresponding breakpoint is enabled in DR7.

Potential Effect on System

Performance counters can not be used for instruction address matches when the corresponding breakpoints are not enabled.

Suggested Workaround

None.

Fix Planned

Yes

162 Writes to Read-Only APIC Register Cause Processor Hang

Description

If a hypervisor (or other software) incorrectly attempts a write to the read-only Extended APIC Feature register (APIC offset 400h) the processor will hang.

Potential Effect on System

A system hang will result if software writes to the read-only Extended APIC Feature register (APIC offset 400h).

Suggested Workaround

Hypervisors (or other software) should not write to the read-only Extended APIC Feature register (APIC offset 400h).

Fix Planned

No

164 DRAM Refresh Controller Not Enabled After Using BIOS Controlled DRAM Initialization

Description

The DRAM controller does not enable the refresh controller when BIOS Controlled DRAM Initialization (Dev 2x7Ch) is used to execute the DRAM initialization sequence described by the JEDEC specification.

Potential Effect on System

Unpredictable system behavior due to memory data loss.

Suggested Workaround

BIOS must use Hardware Controlled DRAM Initialization using DRAM Configuration Low Register[InitDram] (Dev 2x90h[0]); BIOS must not use BIOS Controlled DRAM Initialization.

Fix Planned

Yes

165 #VMEXIT(INVALID) Unconditionally Clears EVENTINJ Field In VMCB

Description

If VMRUN returns with #VMEXIT(INVALID), the EVENTINJ field in the VMCB is unconditionally cleared.

Potential Effect on System

When Guest mode is exited due to VMEXIT(INVALID), the state of the EVENTINJ field in the VMCB is lost and the hypervisor does not have that information available for analysis.

Suggested Workaround

None required. Hypervisors should have this information available in other data structures.

Fix Planned

Yes

166 FXSAVE/FXRSTOR Instructions Use 64-bit Format in Compatibility Mode

Description

In compatibility mode, FXSAVE/FXRSTOR instructions use the 64-bit memory image format (see figure 11-8 in the *AMD64 Architecture Programmer's Manual Volume 2: System Programming*, order# 24593) as opposed to the non-64-bit format (figure 11-9).

Potential Effect on System

A minor performance loss observed due to saving an additional 16 quadwords to memory.

Suggested Workaround

None required.

Fix Planned

Yes

167 System Hang with Chipset Initiated PROCHOT_L Throttling When Clock Divisor is Greater Than 64

Description

When chipset initiated PROCHOT_L throttling is enabled and a clock divisor greater than 64 is selected, a system hang may occur.

Potential Effect on System

The system may hang.

Suggested Workaround

BIOS should program the PMM1 clock divisor to 64 by setting the Power Management Control Low Register [PMM1[ClkSel]] to 011b (Dev 3x80h[14-12] = 011b).

Fix Planned

Yes

168 System May Hang When Exiting C1E or C3

Description

A link may fail to reconnect when exiting C1E or C3.

Potential Effect on System

The link may fail to reconnect when exiting C1E or C3 leading to a system hang.

Suggested Workaround

BIOS should set FidVidEn in the Power Managment Control Low Register[PMM1[FidVidEn]] to 1b (Dev 3x80h[10] = 1b).

When implementing this workaround BIOS must clear AltVidEn in the Power Management Control Low Register[PMM1[AltVidEn]] to 0b (Dev 3x80h[11] = 0b).

Fix Planned

Yes

169 System May Hang Due to DMA or Stalled Probe Response

Description

Under a highly specific and detailed set of internal timing conditions, the Northbridge System Request Queue (SRQ) may stall leading to a deadlock.

Potential Effect on System

Deadlock or machine check exception due to watchdog timer time-out leading to system hang.

Suggested Workaround

The BIOS should set NB_CFG Register[32] (MSR C001_001Fh). No loss of performance results from this workaround.

The workaround for this erratum supersedes the workaround for erratum 131. When implementing this workaround, the workaround for erratum 131 should not be applied.

Fix Planned

No

170 In SVM Mode Incorrect Code Bytes May Be Fetched After A World Switch

Description

On an exit from guest mode (world switch) when CR3 changes, under a highly detailed and specific set of circumstances, incorrect code bytes may be forwarded from the prefetch buffer.

Potential Effect on System

Incorrect code bytes may be executed, resulting in unpredictable system behavior after world switches.

Suggested Workaround

Hypervisors should set TLB_CONTROL to Flush TLB on VMRUN in the VMCB.

Fix Planned

Yes

171 Instruction Break Point On VMRUN Instruction Leads To Unpredictable System Behavior

Description

VMRUN can be interrupted using a hardware instruction breakpoint using one of the debug registers, DR[0-3]. When the debug handler executes IRET, the processor is expected to execute the VMRUN instruction. However, in the failing case, the processor incorrectly re-enters the breakpoint handler with mixed guest and host state. This in turn causes erroneous execution and leads to unpredictable system behavior.

Potential Effect on System

Hypervisor developers will not be able to use hardware instruction break point on VMRUN instruction.

Suggested Workaround

Set the breakpoint on the instruction prior to VMRUN, then single step through VMRUN.

Fix Planned

No

172 Some Registered DIMMs Incompatible With Address Parity Feature

Description

Some DDR-2 registered DIMMs have been observed to generate erroneous address parity errors when operating at 200 MHz MEMCLK . The address parity detection feature may detect false errors upon MEMRESET_L deassertion (during P-state transitions) due to a non-compliant register used on many registered DIMMs already deployed in production.

Potential Effect on System

Machine check exception due to address parity error.

Suggested Workaround

BIOS should disable address parity in the DRAM Configuration Low register (write zero to Dev:2x90[8]).

181 Asserting LDTSTOP_L Before DRAM is Initialized May Cause System Hang

Description

After power up (cold reset), the link frequency is set to 200 MHz, the link width is set to 8, and the processor core frequency may be set below its maximum rated setting. The link frequency and width adjustments can be accomplished using warm resets or LDTSTOP_L. Using LDTSTOP_L for link frequency and width adjustments prior to DRAM initialization may cause the system to hang.

Potential Effect on System

The system may hang.

Suggested Workaround

BIOS can use warm reset to adjust link frequency and width settings instead of LDTSTOP_L. If LDTSTOP_L is used, then BIOS must set DRAM Configuration High[DisDramInterface] to 1 (Dev 2x94h[14] = 1) and set DRAM Configuration High[MemClkFreqVal] to 1 (Dev 2x94h[3] = 1) prior to generating LDTSTOP_L.

Fix Planned

Yes

Documentation Support

The following documents provide additional information regarding the operation of the processor:

- *BIOS and Kernel Developer's Guide for NPT Processors*, order# 32559
- *AMD64 Architecture Programmer's Manual Volume 1: Application Programming*, order# 24592
- *AMD64 Architecture Programmer's Manual Volume 2: System Programming*, order# 24593
- *AMD64 Architecture Programmer's Manual Volume 3: General-Purpose and System Instructions*, order# 24594
- *AMD64 Architecture Programmer's Manual Volume 4: 128-Bit Media Instructions*, order# 26568
- *AMD64 Architecture Programmer's Manual Volume 5: 64-Bit Media and x87 Floating-Point Instructions*, order# 26569
- *AMD CPUID Specification*, order# 25481.

See the AMD Web site at www.amd.com for the latest updates to documents. For documents subject to a non-disclosure agreement (NDA), please contact your local sales representative.